

WHAT IS CLAIMED IS:

1. A primary controller for receiving and executing database query and write commands in a distributed database system, the primary controller comprising:

a table distribution processing module for receiving a database write command and determining a node on which data for said database write command is to be stored;

a pivot table processing module for determining a pivot table from a plurality of join tables;

a query distribution processing module for sending a database query command, determining how to build a join table, and sending information on how to build said join table;

a results file receipt processing module for receiving an intermediate results file and storing said intermediate results file; and

a final results file build processing module for processing said intermediate results file and building a final results file that represents the final results of said database query command.

2. The primary controller of Claim 1, further comprising a results file post processing module for performing post-processing operations on said final results file.

3. The primary controller of Claim 2, wherein said post processing operations comprise removing duplicate entries from said final results file.

4. The primary controller of Claim 1, further comprising a node interface processing module for sending outgoing data from the primary front end processor and receiving incoming data for the primary front end processor.

5. A primary controller for receiving and executing a database query command and write command in a distributed database system having a plurality of database tables stored on more than one node, a portion of each database table stored on each node, the primary controller comprising:

a node interface processing module configured to receive a primary query command and a write command;

a table distribution processing module configured to receive said write command from said node interface processing module and determine a node on which data for said database write command is to be stored;

a pivot table processing module configured to receive said primary query command from said node interface processing module and determine which of said plurality of database tables is to be a pivot table, wherein said pivot table is selected from said plurality of databases tables and is not sent to other nodes;

a query distribution processing module configured to  
send a secondary query command to the other nodes,  
determine a join table build definition, and  
send information on how to build said join table to the other nodes;

a results file receipt processing module for receiving an intermediate results file and storing said intermediate results file; and

a final results file build processing module for processing said intermediate results file and building a final results file that represents the final results of said primary query command.

6. The primary controller of Claim 5, further comprising a results file post processing module for performing post-processing operations on said final results file.

7. The primary controller of Claim 6, wherein said post processing operations comprise removing duplicate entries from said final results file.

8. The primary controller of Claim 7, wherein said post processing operations further comprise ordering entries in said final results file.

9. The primary controller of Claim 5, further comprising a node interface processing module for sending outgoing data from the primary front end processor and receiving incoming data for the primary front end processor.

10. A method of receiving and executing database query and write commands in a distributed database system, the method comprising:

receiving a database write command and determining a node on which data for said database write command is to be stored;  
determining a pivot table from a plurality of join tables;  
sending a database query command;

determining how to build a join table;  
sending information on how to build said join table;  
receiving an intermediate results file and storing said intermediate results file;  
and

processing said intermediate results file and building a final results file that represents the final results of said database query command.

11. The method of Claim 10, further comprising performing post-processing operations on said final results file.

12. The method of Claim 11, wherein said post processing operations comprise removing duplicate entries from said final results file.

13. A method of receiving and executing a database query command and write command in a distributed database system having a plurality of database tables stored on more than one node, a portion of each database table stored on each node, the method comprising:

receiving a primary query command and a write command;  
receiving said write command and determining a node on which data for said database write command is to be stored;  
receiving said primary query command and determining which of said plurality of database tables is to be a pivot table, wherein said pivot table is selected from said plurality of databases tables and is not sent to other nodes;  
sending a secondary query command to the other nodes;  
determining a join table build definition;  
sending information on how to build said join table to the other nodes;  
receiving an intermediate results file and storing said intermediate results file;  
and

processing said intermediate results file and building a final results file that represents the final results of said primary query command.

14. The method Claim 13, further comprising performing post-processing operations on said final results file.

15. The primary controller of Claim 14, wherein said post processing operations comprise removing duplicate entries from said final results file.

16. The primary controller of Claim 14, wherein said post processing operations comprise ordering entries in said final results file.